arXiv:1512.03531v4 [cs.CC] 17 Jun 2016

Constructive non-commutative rank computation in deterministic polynomial time over fields of arbitrary characteristics

An addendum to the paper entitled

Non-commutative Edmonds' problem and matrix semi-invariants

Gábor Ivanyos* Youming Qiao† K. V. Subrahmanyam ‡

June 20, 2016

Abstract

We extend our techniques developed in our manuscript mentioned in the subtitle to obtain a deterministic polynomial-time algorithm for computing the non-commutative rank together with certificates of linear spaces of matrices over sufficiently large base fields.

The key new idea is a reduction procedure that keeps the blow-up parameter small, and there are two methods to implement this idea: the first one is a greedy argument that removes certain rows and columns, and the second one is an efficient algorithmic version of a result of Derksen and Makam. Both methods rely crucially on the regularity lemma in the aforementioned manuscript, and in this note we improve that lemma by removing a coprime condition there.

2010 Mathematics Subject Classification: Primary 13A50, 68W30.

Keywords: Edmonds' problem, symbolic determinant identity test, semi-invariants of quivers, non-commutative rank

1 Introduction

Let $\mathbb{F}$ be a field. Let $M(n\times m,\mathbb{F})$ be the linear space of $n\times m$ matrices over $\mathbb{F}$; if $n=m$ we may just write as $M(n,\mathbb{F})$. We call a linear subspace of $M(n\times m,\mathbb{F})$ a *matrix space*. Then, given a matrix space $\mathcal{B}\leq M(n,\mathbb{F})$, its (k,ℓ)-*blow-up* $\mathcal{B}^{\{k,\ell\}}$ is defined as the matrix space $\mathcal{B}\otimes M(k\times\ell,\mathbb{F})$ in $M(nk\times n\ell,\mathbb{F})$. Although the most relevant blow-ups in this context are square (e.g, of the form $\mathcal{B}^{\{k,k\}}$), non-square blow-ups turned out to be crucial in the reduction techniques in [DM15].

From earlier results, we can define the *non-commutative rank* of $\mathcal{B}\leq M(n,\mathbb{F})$ to be the maximum over d of $\frac{1}{d}$ times the maximum rank of a matrix from the blow-up $\mathcal{B}^{\{d,d\}}$. An important question is to determine bounds on the blow-up parameter d (as a function of n) which achieves the desired maximum. From the work of [Der01] it is known that $d=O(n^4\cdot 4^{n^2})$, over algebraically closed fields of characteristic zero.[1] In [IQS15] it is shown that $d=O(n^4\cdot 2^{O(n\log n)})$ over large fields of arbitrary characteristics.

These exponential bounds turned out to be sufficient in [GGOW15] to compute the non-commutative rank in deterministic polynomial time, over fields of characteristic zero, using an algorithm of Gurvits [Gur04]. After [GGOW15], the following problems were still open: existence of

*Institute for Computer Science and Control, Hungarian Academy of Sciences, Budapest, Hungary. E-mail: Gabor.Ivanyos@sztaki.mta.hu.

†Centre for Quantum Computation and Intelligent Systems, University of Technology, Sydney, Australia. E-mail: jimmyqiao86@gmail.com.

‡Chennai Mathematical Institute, Chennai, India. E-mail: kv@cmi.ac.in.

[1] Derksen's result applies to a wide class of invariant rings.

(1) a polynomial-time algorithm for the problem over finite fields, and

(2) a search version of the problem, that is, explicitly exhibiting a matrix of rank rd in the d-th blow-up and a proof that the non-commutative rank is at most r, even over fields of characteristic 0.

Recently, Derksen and Makam proved in [DM15] that it suffices to take the maximum over d between 1 and $n-1$ for sufficiently large fields, by discovering a concavity property of blow-ups, and using the regularity lemma of blow-ups from our previous manuscript [IQS15]. In the first version of this note, by showing that the concavity property can be made constructive, and building on the techniques from [IQS15], we obtained a deterministic polynomial-time algorithm for the non-commutative rank problem, which is constructive and works over large enough fields regardless of the characteristics. This answers the two open problems just mentioned. After the first version of this note appeared on the arXiv, we discovered that a very simple observation already gives us the result, without having to use the results from Derksen and Makam. This argument also gives a different proof that the nullcone of the matrix semi-invariants (for this concept see [IQS15, DM15]) is defined by those invariants of degree no more than $O(n^2)$.

To present our main theorem formally, let us recall some basic facts about the non-commutative rank. A subspace V of $\mathbb{F}^n$ is an *s-shrunk subspace* of $\mathcal{B}$ if $\dim \mathcal{B}(V) \leq \dim V - s$. If $\mathcal{B}$ has an s-shrunk subspace, then no matrix from $\mathcal{B}$ can have rank larger than $n-s$. As by a square d-blow-up, s-shrunk subspaces are blown up to ds-shrunk subspaces, the existence of an s-shrunk subspace implies an upper bound $n-s$ for the non-commutative rank of $\mathcal{B}$. Fortin and Reutenauer showed in [FR04] that the minimum of the number $n-s$ such that there exist s-shrunk subspaces coincides with the non-commutative rank.

Our main theorem is then as follows.

Theorem 1.1. *Let $\mathcal{B} \leq M(n, \mathbb{F})$ be a matrix space given by a linear basis, and suppose $|\mathbb{F}| = n^{\Omega(1)}$. Suppose that $\mathcal{B}$ has (a priori unknown) non-commutative rank r. Then there is a deterministic algorithm using $n^{O(1)}$ arithmetic operations over $\mathbb{F}$ that constructs a matrix of rank rd in a blow-up $\mathcal{B}^{\{d,d\}}$ for some $d \leq r+1$ as well as an $(n-r)$-shrunk subspace of $\mathbb{F}^n$ for $\mathcal{B}$. When $\mathbb{F} = \mathbb{Q}$, the final data as well as all the intermediate data have size polynomial in the size of the input data and hence the algorithm runs in polynomial time.*

Remark 1.2. (a) If the method of Derksen and Makam [DM15] is used as in Section 4, then in the above theorem we can improve the parameter slightly to $d \leq r-1$ instead of $d \leq r+1$.

(b) Polynomiality of the algorithm can also be proved for a wide range "concrete" base fields $\mathbb{F}$. These include sufficiently large finite fields, and also number fields and transcendental extensions of constant degree over finite fields and over number fields.

(c) In particular, the non-commutative rank can be computed in deterministic polynomial time in positive characteristic as well, assuming that the ground field is sufficiently large.

Our result also settles a question of Gurvits [Gur04], asking if it is possible to decide efficiently, over fields of positive characteristics, whether or not there exists a nonsingular matrix in a matrix space having the so called Edmonds-Rado property. A matrix space has the *Edmonds-Rado property* if it satisfies the promise that it either contains a nonsingular matrix, or it shrinks some subspace.

Since the algorithm in Theorem 1.1 efficiently tells whether the given matrix space has a shrunk subspace (e.g. the non-commutative rank is not full), it settles Gurvits' question, when the field size is as stated in the hypothesis. However we can say more, even when the base field is a "too small" finite field. We can first apply Theorem 1.1 to a sufficiently large extension. We can embed the extension field as a maximal commutative subfield of the full matrix algebra of an appropriate degree over the original field, using the regular representation. We obtain the following, therefore settling Gurvits' question over any field.

Corollary 1.3. *Let $\mathcal{B} \leq M(n, \mathbb{F})$ be a matrix space given by a basis, where $\mathbb{F}$ is a finite field. Suppose that $\mathcal{B}$ has (a priori unknown) non-commutative rank r. Then there is a deterministic polynomial time algorithm that constructs a matrix of rank rd in a blow-up $\mathcal{B}^{\{d,d\}}$ for some $d \leq O(r \log_{|\mathbb{F}|} r)$ as well as an $(n-r)$-shrunk subspace of $\mathbb{F}^n$ for $\mathcal{B}$.*

Techniques. Theorem 1.1 is built on our previous techniques in [IQS15], and a key new idea. The algorithm in [IQS15] takes exponential time, because we increase the blow-up size in an iterative way, and in each iteration the blow-up size is increased multiplicatively by the "scaled" rank. The key new insight of Derksen and Makam is that we can keep the blow-up size small: when the scaled rank is r, then the blow-up size can be brought back to $O(r)$! As mentioned, we offer two methods to realize this reduction idea: a simpler method from us, and a method based on the technique of Derksen and Makam [DM15].

We also provide a technical improvement to a critical lemma in our previous paper, namely the constructive regularity lemma. Mathematically, that lemma says that, when $\mathbb{F}$ is large enough, then $\mathrm{rk}(\mathcal{B}\otimes M(d,\mathbb{F}))$ is divisible by d [IQS15, Lemma 5.6]. We use it in the algorithm in the following situation: given $A \in \mathcal{B}\otimes M(d,\mathbb{F})$ of rank $(r-1)d+k$ where $1<k<d$, we want to construct $A' \in \mathcal{B}\otimes M(d,\mathbb{F})$ of rank $\geq rd$ efficiently. In [IQS15, Lemma 5.7], this was achieved under the condition that, if $\mathrm{char}(\mathbb{F})=p>0$, then $p \nmid d$. In this note, we remove this coprime condition.

Organization. We will prove the full regularity lemma first in Section 2. We then prove the main theorem, Theorem 1.1 in Section 3, using our blow-up reduction method. Finally in Section 4 we show that the Derksen–Makam technique can be constructivized to provide another blow-up reduction method.

2 The full regularity lemma

We first present the formal statement of the regularity lemma in its full generality. We also add a technical notion that will be useful for the proof of Theorem 1.1. Let $n \in \mathbb{N}$, and let $\mathbf{i}=(i_1,\ldots,i_r)$, $\mathbf{j}=(j_1,\ldots,j_r)$ be two sequences of integers, where $1 \leq i_1 < \cdots < i_r \leq n$ and $1 \leq j_1 < \cdots < j_r \leq n$. For a matrix $A \in M(n,\mathbb{F})\otimes M(d,\mathbb{F})$, the $r\times r$ *window* indexed by $\mathbf{i}$, $\mathbf{j}$ is the sub-matrix of A consisting of the blocks indexed by (i_k,j_ℓ), $k,\ell \in [r]$.

Lemma 2.1 (Regularity of blow-ups)**.** *For $\mathcal{B} \leq M(n,\mathbb{F})$ and $\mathcal{A}=\mathcal{B}^{\{d,d\}}$, assume that $|\mathbb{F}|=(rd)^{\Omega(1)}$. Given a matrix $A\in\mathcal{A}$ with $\mathrm{rk}A>(r-1)d$, there exists a deterministic algorithm that returns $\widetilde{A}\in\mathcal{A}$ and an $r\times r$ window W in $\widetilde{A}$ such that W is nonsingular (of rank rd). This algorithm uses $\mathrm{poly}(nd)$ arithmetic operations and, over $\mathbb{Q}$, the algorithm runs in polynomial time. In particular, all intermediate numbers have bit lengths polynomial in the input size.*

The cases (a) $\mathrm{char}(\mathbb{F})=0$, (b) $\mathrm{char}(\mathbb{F})$ and d are coprime, and $|\mathbb{F}|=(rd)^{\Omega(1)}$ were settled in [IQS15, Lemma 5.7]. The main issue with the case when d is divisible by $\mathrm{char}(\mathbb{F})$ was that then the proof would have required an efficient construction of an appropriate Artin-Schreier-Witt extension of $\mathbb{F}_p(x)$, which was unavailable to us at the time of writing [IQS15]. Now we have such a construction in Lemma 2.4.

The proof makes use of the following two results from [IQS15].

Proposition 2.2 ([IQS15, Proposition 4.4])**.** *Let $\mathbb{L}$ be a cyclic extension of degree d of a field $\mathbb{K}$, and suppose that $\mathbb{L}$ is given by structure constants w.r.t. a $\mathbb{K}$-basis $A_1,\ldots,A_d$. Similarly, a generator σ for the Galois group is assumed to be given by its matrix in terms of the same basis. Let Y be a formal variable. Then one can construct a $\mathbb{K}(Y)$-basis Γ of $M(d,\mathbb{K}(Y))$ such that the $\mathbb{K}(Y^d)$-linear span of Γ is a central division algebra over $\mathbb{K}(Y^d)$ of index d, using $\mathrm{poly}(d)$ arithmetic operations in $\mathbb{K}$. Furthermore for $\mathbb{K}=\mathbb{Q}[\sqrt[d]{1}]$, the bit complexity of the algorithm (as well as the size of the output) is also $\mathrm{poly}(d)$.*

Lemma 2.3 (Conditional regularity [IQS15, Lemma 5.4])**.** *Assume that we are given a matrix $A\in\mathcal{B}^{\{d\}}\leq M(dn,\mathbb{F})$ with $\mathrm{rk}(A)=(r-1)d+k$ for some $1<k<d$. Let X and Y be formal variables and put $\mathbb{K}=\mathbb{F}'(X)$, where $\mathbb{F}'$ is a finite extension of $\mathbb{F}$ of degree at most d. Suppose further that $|\mathbb{F}|>(nd)^{O(1)}$ and that we are also given a $\mathbb{K}(Y)$-basis Γ of $M(d,\mathbb{K}(Y))$ such that the $\mathbb{K}(Y^d)$-linear span of Γ is a central division algebra D' over $\mathbb{K}(Y^d)$. Let δ be the maximum of the degrees of the polynomials appearing as numerators or denominators of the entries of the matrices in Γ. Then, using $(nd+\delta)^{O(1)}$ arithmetic operations in $\mathbb{F}$, one can find a matrix $A''\in\mathcal{B}^{\{d\}}$ with $\mathrm{rk}(A'')\geq rd$. Furthermore, over $\mathbb{Q}$ the bit complexity of the algorithm is polynomial in the size of the input data (that is, the total number of bits describing the entries of matrices and in the coefficients of polynomials).*

Proof of Lemma 2.1. The statement, except the window part, readily follows by plugging Lemma 2.5 of the next subsection and Proposition 2.2 to Lemma 2.3. To see that such a window can be computed, we first observe that the lemma applies to d-blow-ups of rectangular matrices, by simple zero padding. Second, apply the lemma and find an $rd \times rd$ nonsingular submatrix of the given matrix A. If the column indices include some such that not all of its $d-1$ siblings are included, then (1) delete the corresponding column from the original matrix space; (2) let A' be the matrix obtained by deleting the corresponding d columns from A. Then $\mathrm{rk}(A') > \mathrm{rk}(A) - (d-1)$. So we apply the regularity lemma in the rectangular space with A', to round up the rank to $\mathrm{rk}(A)$ again. Do the same for row indices. Iterate until we obtain a full window. □

2.1 Efficient construction of cyclic field extensions of arbitrary degrees

A cyclic extension of a field $\mathbb{K}$ is a finite Galois extension of $\mathbb{K}$ having a cyclic Galois group. By constructing a cyclic extension $\mathbb{L}$ we mean constructing the extension as an algebra over $\mathbb{K}$, e.g., by giving an array of *structure constants* with respect to a $\mathbb{K}$-basis for $\mathbb{L}$ defining the multiplication on $\mathbb{L}$ as well as specifying a generator of the Galois group, e.g, by its matrix with respect to a $\mathbb{K}$-basis.

Lemma 2.4. *Given a prime p and an integer $s \geq 1$, one can construct in time $\mathrm{poly}(p^s)$ a cyclic extension K_s of $\mathbb{F}_p(Z)$ of degree p^s such that $\mathbb{F}_p$ is algebraically closed in K_s. The field K_s will be given in terms of structure constants with respect to a basis over $\mathbb{F}_p(Z)$, and the generator σ for the Galois group will be given by its matrix in terms of the same basis. The structure constants as well as the entries of the matrix for σ will be polynomials in $\mathbb{F}_p[Z]$ of degree $\mathrm{poly}(p^s)$.*

Proof. First we briefly recall the general construction given in Section 6.4 of [Ram54]. This, starting from a field K_0 of characteristic p, recursively builds a tower $K_0 < K_1 < \ldots < K_s$ of fields such that K_j is a cyclic extension of K_0 of degree p^j. Assume that K_s together with a K_0-automorphism σ_s of order p^s has already been constructed. (Initially let σ_0 be the identity map on K_0.) Then for any element $\beta_s \in K_s$ with $\mathrm{Tr}_{K_s:K_0}(\beta_s) = 1$ and for any $\alpha_s \in K_s$ such that $\alpha_s^{\sigma_s} - \alpha_s = \beta_s^p - \beta_s$ the polynomial $X^p - X - \alpha_s$ is irreducible in $K_s[X]$. (Existence of α_s with the required property follows from the additive Hilbert 90.) Put $K_{s+1} = K_s[X]/(X^p - X - \alpha_s)$ and let $\omega_{s+1} \in K_{s+1}$ be the image of of X under the projection $K_s[X] \to K_{s+1}$. Then σ_s extends to a K_0-automorphism σ_{s+1} of degree p^{s+1} of K_{s+1} such that $\omega_{s+1}^{\sigma_{s+1}} = \omega_{s+1} + \beta_s$. This gives a cyclic extension of degree p^{s+1}.

Now we specify some details of a polynomial time construction for $K_0 = \mathbb{F}_p(Z)$ following the method outlined above. In the first step we take $\beta_0 = 1$, and, in order to guarantee that the only elements in K_1 which are algebraic over $\mathbb{F}_p$ is F_p (we also use the phrase F_p is algebraically closed in K_1 when this property holds), we take $\alpha_0 = Z$. Then K_1 is a pure transcendental extension of $\mathbb{F}_p$. As K_s/K_0 is a cyclic extension of oder p^s, it has a unique subfield which is an order p extension of K_0. This must be K_1. Then $\mathbb{F}_p$ has no proper finite extension in K_s as otherwise K_0 would also have another degree p extension.

We consider the following K_0-basis for K_s:

$$\Gamma_s = \left\{ \prod_{j=1}^{s} \omega_j^k, \quad (k = 0, \ldots, p-1) \right\},$$

where ω_j is a root of $X^p - X - \alpha_{j-1}$ in K_j. We claim that $\mathrm{Tr}_{K_j:K_{j-1}}(\omega_j^{p-1}) = -1$. Indeed, in the K_{j-1}-basis $\omega_j^0, \ldots, \omega_j^{p-1}$ for K_j, in the matrix of multiplication by ω_j^{p-1} the diagonal entries consist of $p-1$ ones and one zero. Therefore $\mathrm{Tr}_{K_j:K_{j-1}}(\omega_j^{p-1}\gamma) = -\gamma$ for every $\gamma \in K_{j-1}$, whence $\mathrm{Tr}_{K_j:K_0}(\omega_j^{p-1}\gamma) = -\mathrm{Tr}_{K_{j-1}:K_0}(\gamma)$. Now by induction we obtain $\mathrm{Tr}_{K_j:K_0} \prod_{i=1}^{j} \omega_i^{p-1} = (-1)^j$. Therefore in each step (when $j > 0$) we can choose $\beta_j = (-1)^j \prod_{i=1}^{j} \omega_i^{p-1}$ and α_j thereafter, following the construction in the standard proof of the additive Hilbert 90. Specifically, we set

$$\alpha_j = (-1)^{j+1} \sum_{k=1}^{p^j-1} \beta_j^{\sigma_j^k} \left(\sum_{\ell=0}^{k-1} (\beta_j^p - \beta_j)^{\sigma_j^\ell} \right). \tag{2.1}$$

Then $\alpha_j^{\sigma_j} - \alpha_j = \beta_j^p - \beta_j$. Notice that α_j is a sum of terms with each of which, up to a sign, is a product of at most $p+1$ conjugates $\beta_j^{\sigma_j^\ell}$ (with various ℓs) of β_j ($\ell \le p^j$)

Assume by induction that the structure constants of K_j with respect to the basis Γ_j are polynomials from $\mathbb{F}_p[Z]$ of degree at most Δ_j and the same holds for the entries of the matrix of σ_j^ℓ for every $1 \le \ell < p^j$ (written in the same basis). For $j = 1$ this holds with $\Delta_1 = 1$. (To see this, observe that for $0 \le k, \ell < p$, the product $\omega_1^k \omega_1^\ell$ is the basis element of $\omega_1^{k+\ell}$ if $k + \ell < p$, while otherwise it equals the sum $\omega_1^{k+\ell-p+1} + Z\omega_1^{k+\ell-p}$.) Then, if we express α_j in terms of the basis Γ_j using Eq. 2.1, we obtain that its coordinates are polynomials of degree at most $(2p+1)\Delta_j$. This is because $(-1)^j \beta_j \in \Gamma_j$, whence $\beta_j^{\sigma^\ell}$ has coordinates of polynomials of degree bounded by Δ_j. In Eq. 2.1, we have the products of at most $p+1$ such elements, so the result will have polynomial coordinates of degree at most $(2p+1)\Delta_j$.

Now consider the product of two elements $\omega_{j+1}^k \gamma_1$ and $\omega_{j+1}^\ell \gamma_2$ of Γ_{j+1}. Here $k, \ell < p$ and $\gamma_1, \gamma_2 \in \Gamma_j$. The coordinates of the product $\gamma_1\gamma_2$ with respect to Γ_j are polynomials of degree at most Δ_j. The same holds for the product $\omega_{j+1}^{k+\ell}\gamma_1\gamma_2$ if $k + \ell < p$. If $k + \ell > p$, then $\omega_{j+1}^{k+\ell} = \omega_{j+1}^p \omega_{j+1}^{k+\ell-p} = (\omega_{j+1} + \alpha_j)\omega_{j+1}^{k+\ell-p}$, whence $\omega_{j+1}^{k+\ell}\gamma_1\gamma_2$ is the sum of $\omega_{j+1}^{1+k+\ell-p}\gamma_1\gamma_2$ and $\alpha_j\gamma_1\gamma_2$. The former term has coordinates of degree at most Δ_j, the coordinates of the latter are polynomials of degree at most $(2p+1)\Delta_j + \Delta_j + \Delta_j = (2p+3)\Delta_j$.

Now consider the conjugate of $\omega_{j+1}^k \gamma$ by σ_{j+1}^ℓ, where $1 \le \ell < p^{j+1}$, $1 \le k \le p-1$ and $\gamma \in \Gamma_j$. This conjugate is $(\omega_{j+1}^{\sigma_{j+1}^\ell})^k \gamma^{\sigma_{j+1}^\ell}$. The second term equals $\gamma^{\sigma_j^\ell}$ which has coordinates of degree at most Δ_j. To investigate the first term, recall that $\omega_{j+1}^{\sigma_{j+1}} = \omega_{j+1} + \beta_j$, whence

$$\omega_{j+1}^{\sigma_{j+1}^\ell} = \omega_{j+1} + \sum_{r=0}^{\ell-1} \beta_j^{\sigma_j^r}$$

The element $\delta = \sum_{r=0}^{\ell-1} \beta_j^{\sigma_j^r}$, expressed in terms of Γ_j, has again polynomial coordinates of degree at most Δ_j. Then $(\omega_{j+1}^{\sigma_{j+1}^\ell})^k$ is the sum (with binomial coefficients) of terms of the form $\omega_{j+1}^r \delta^{k-r}$. The power δ^{k-r} has coordinates of degree at most $(k-r)\Delta_j + (k-r-1)\Delta_j \le (2p-1)\Delta_j$ in terms of Γ_j, whence we conclude that $(\omega_{j+1}^{\sigma_{j+1}^\ell})^k$ has, in terms of Γ_{j+1} polynomial coordinates of degree at most $(2p-1)\Delta_j$. It follows that the matrix of any power of σ_{j+1} has polynomial entries of degree at most $2p\Delta_j$.

We obtained that the function $(2p+3)^s = \mathrm{poly}(p^s)$ is an upper bound for both the structure constants and for the matrices of the powers of σ_s. □

Lemma 2.5. *Let $\mathbb{F}'$ be a field. Let d be any non-negative integer. If* $\mathrm{char}(\mathbb{F}') = 0$ *then $d_1 = d$. If* $\mathrm{char}(\mathbb{F}') = p > 0$ *then let d_1 be the p-free part of d, that is, $d = d_1 p^s$, where $p \nmid d_1$ and $s \in \mathbb{N}$. Assume that $\mathbb{F}'$ contains a known d_1th root of unity ζ. Then a cyclic extension $\mathbb{L}$ degree d of $\mathbb{K} := \mathbb{F}'(X)$ can be computed using* $\mathrm{poly}(d)$ *arithmetic operations. $\mathbb{L}$ will be given by structure constants with respect to a basis, and the matrix for a generator of the Galois group in terms of the same basis will also be given. All the output entries (the structure constants as well as the entries of the matrix representing the Galois group generator) will be polynomials of degree* $\mathrm{poly}(d)$ *in $\mathbb{F}'[X]$. Furthermore for $\mathbb{F}' = \mathbb{Q}[\sqrt[d_1]{1}]$, the bit complexity of the algorithm (as well as the size of the output) is* $\mathrm{poly}(d)$.

Proof. Put $\mathbb{L}_1 = \mathbb{F}'(Y)$ and $X = Y_1^{d_1}$. Then $1, Y_1, \ldots, Y_1^{d_1}$ are a $\mathbb{F}'(X)$-basis for $\mathbb{L}_1$ with $Y_1^i Y_1^j = Y_1^{i+j}$ if $i + j \le d_1$ and $XY_1^{i+j-d_1}$ otherwise. Further note that the linear extension σ_1 of the map sending Y_1^j to $\zeta^j Y_1^j$ is an automorphism of degree d_1. Then $\mathbb{L}_1$ is a cyclic extension of $\mathbb{F}'(X)$ of degree d_1. This procedure has been used in [IQS15].

We can compute whether $\mathrm{char}(\mathbb{F}')$ is a divisor of d by testing the multiples of the identity element up to d. If $\mathrm{char}(\mathbb{F}') = 0$, or if $\mathrm{char}(\mathbb{F}') = p > 0$ and $p \nmid d$, we are done. Note that in the following $p \le d$.

If $\mathrm{char}(\mathbb{F}') = p > 0$ and $p \mid d$, let d_1 be in the statement, so $d = d_1 p^s$. Let $d_2 = p^s$, and $\mathbb{F}_p$ be the prime field of $\mathbb{F}'$. Construct the cyclic extension of degree d_2 of $\mathbb{F}_p(X)$ over $\mathbb{F}_p$ by Lemma 2.4, and let the resulting field be $\mathbb{L}_2$. We also obtain the matrix a generator σ_2 of the Galois group. Then put $\mathbb{L} = \mathbb{L}_1 \otimes_{\mathbb{F}_p(X)} \mathbb{L}_2$. It contains a copy of $\mathbb{K} = \mathbb{F}'(X) \cong \mathbb{F}'(X) \otimes_{\mathbb{F}_p(X)} \mathbb{F}_p(X)$. We take the product basis for the structure constants and for matrix representation of the automorphism $\sigma_1 \otimes \sigma_2$. □

3 Proof of the main theorem

We recall some preparation material from [IQS15].

Finding an sd-shrunk subspace for the $\mathcal{B}^{\{d,d\}}$ is equivalent to finding an s-shrunk subspace for $\mathcal{B}$ because of the following simple observations ([IQS15, Proposition 5.2]). Firstly, for every s-shrunk subspace U of $\mathbb{F}^n$ the subspace $U \otimes \mathbb{F}^d$ for $\mathcal{B}$ is obviously an sd-shrunk subspace for $\mathcal{B}^{\{d,d\}}$. Conversely, a s'-shrunk subspace for $\mathcal{B}^{\{d,d\}}$ can be easily embedded into a subspace of the form $U \otimes \mathbb{F}^d$ where U is an s-shrunk subspace for $\mathcal{B}$ with $sd \geq s'$.

We shall also need the following useful lemma.

Lemma 3.1 (Data reduction, [dGIR96] and [IQS15, Lemma 5.3])**.** *Let $\mathcal{B} \leq M(k \times \ell, \mathbb{F})$ be given by a basis $B_1, \ldots, B_m$, and let $\mathbb{K}$ be an extension field of $\mathbb{F}$. Let S be a subset of $\mathbb{F}$ of size at least $r+1$. Suppose that we are given a matrix $A' = \sum a'_i B_i \in \mathcal{B} \otimes_{\mathbb{F}} \mathbb{K}$ of rank at least r. Then we can find $A = \sum a_i B_i$ of rank also at least r with $a_i \in S$. The algorithm uses* $\mathrm{poly}(k, \ell, r)$ *rank computations for matrices of the form $\sum \gamma_i A_i$ where $a''_i \in \{a'_1, \ldots, a'_m\} \cup S$.*

The main technical ingredient of our algorithm will be the following result from [IQS15] (Theorem 5.10 of ibid.). It states that either a shrunk subspace witnessing that the (scaled down) rank of a matrix in a blow-up reaches the non-commutative rank or a matrix in a larger blow-up having larger scaled down rank can be efficiently constructed.

Theorem 3.2 ([IQS15, Theorem 5.10])**.** *Let $\mathcal{B} \leq M(n, \mathbb{F})$ and let $\mathcal{A} = \mathcal{B}^{\{d,d\}}$. Assume that we are given a matrix $A \in \mathcal{A}$ with $\mathrm{rk}(A) = rd$, and $|\mathbb{F}|$ is $(ndd')^{\Omega(1)}$, where $d' > r$ is any positive integer. There exists a deterministic algorithm that returns either an $(n-r)d$-shrunk subspace for $\mathcal{A}$ (equivalently, an $(n-r)$-shrunk subspace for $\mathcal{B}$), or a matrix $B \in \mathcal{A} \otimes M(d', \mathbb{F})$ of rank at least $(r+1)dd'$. To be more specific, in this case an $(r+1) \times (r+1)$ window is also found such that the corresponding $(r+1)dd'$ times $(r+1)dd'$ sub-matrix of B has full rank. This algorithm uses* $\mathrm{poly}(ndd')$ *arithmetic operations and, over $\mathbb{Q}$, all intermediate numbers have bit lengths polynomial in the input size.*

The sentence on the $(r+1) \times (r+1)$ window is not explicitly stated in [IQS15]. However, the algorithm in its proof contains, as a last step, a call to the method behind Lemma 2.1. Also, the theorem was stated only under the assumption that d was not divisible by $\mathrm{char}(\mathbb{F})$ because of this last call. As the algorithm up to this step constructs a matrix of rank greater than rdd', the complete version regularity, as stated in Lemma 2.1, makes it possible to dispense with that assumption.

To obtain the algorithm for Theorem 1.1, the regularity lemma needs to be accompanied with a reduction procedure that keeps the blow-up parameter small. We mentioned in the introduction that there are two methods for this purpose, and in this section we use our method. The method based on the Derksen–Makam technique will be presented in Section 4.

Lemma 3.3. *Let $\mathcal{B} \leq M(n, \mathbb{F})$, and $d > n+1$. Assume we are given a matrix $A \in \mathcal{B}^{\{d,d\}}$ of rank dn. Then there exists a deterministic polynomial-time procedure that constructs $A' \in \mathcal{B}^{\{d-1,d-1\}}$ of rank $(d-1)n$.*

Proof. Let A'' be an appropriate $(d-1)n \times (d-1)n$ submatrix of A corresponding a matrix in $\mathcal{B}^{\{d-1,d-1\}}$. We claim A'' is of rank $> (d-1)(n-1)$. Suppose not, as A is obtained from A'' from adding n rows and then n columns, and $d > n+1$, we have $\mathrm{rk}(A) \leq \mathrm{rk}(A'') + 2n \leq dn - d - n + 1 + 2n < dn$, a contradiction. Now that $\mathrm{rk}(A'') > (d-1)(n-1)$, using Lemma 2.1, we obtain $A' \leq \mathcal{B}^{\{d-1,d-1\}}$ of rank $(d-1)n$. □

Proof of Theorem 1.1. Let $B_1, \ldots, B_m$ be the input basis for $\mathcal{B}$. The algorithm is an iteration based on Theorem 3.2. During the iteration we have a matrix $A = \sum_i B_i \otimes T_i \in \mathcal{B}^{\{d,d\}}$ of rank rd for some integer $d \leq r-1$. We assume that at least one of the basis elements B_i has rank larger than 1, as otherwise [IKS10] works even with $d = 1$. Initially $d = 1$ and A is a basis element having rank at least 2. The procedure behind Theorem 3.2 either returns an $(n-r)$-shrunk subspace (in which case we are done), or a new matrix (denoted also by A) in a blow-up $\mathcal{B}^{\{d',d'\}}$ of rank at least $(r+1)d'$ for some $d' < r^2$ together with a square window of size $r+1$ so that the corresponding sub-matrix of A is of rank $(r+1)d'$. If $d' > r+2$ then we

apply Lemma 3.3. Note that n in the statement will be $r+1$, and we shall use it repeatedly to get a matrix in the $(r+2, r+2)$-blow-up with a similar content as above. The main content of this are $r+2$ by $r+2$ matrices $T_1, \ldots, T_m$ such that the $(r+1)(r+2) \times (r+1)(r+2)$ sub-matrix of $A' = \sum B_i \otimes T_i$ has full rank. Then we replace A with A' and apply the size reduction procedure in Lemma 3.1 to arrange that the entries of T_i fall into the prescribed subset of $\mathbb{F}$, and continue the iteration with this new matrix A. □

4 Constructivizing the result of Derksen and Makam

Here is an algorithmic version of Lemma 2.7 of [DM15].

Lemma 4.1. *Let $\mathcal{B} \leq M(n, \mathbb{F})$. Assume that for $k, \ell = 1, \ldots, N$ we are given matrices $M_0(k, \ell) \in \mathcal{B}^{\{k,\ell\}}$ of rank $r_0(k, \ell)$, and suppose that $|\mathbb{F}| \geq 2nN + 1$. Then for every $k, \ell = 0, \ldots, N$ we can efficiently (that is, by an algorithm that uses* poly(Nn) *arithmetic operations and, over e.g. $\mathbb{Q}$, produces intermediate and final data of size polynomial in the input size) construct matrices $M(k, \ell) \in \mathcal{B}^{\{k,\ell\}}$ of rank $r(k, \ell) \geq r_0(k, \ell)$ such that*

(1) $r(k, \ell+1) \geq r(k, \ell)$ $(0 \leq \ell < N)$;

(2) $r(k+1, \ell) \geq r(k, \ell)$ $(0 \leq k < N)$;

(3) $r(k, \ell+1) \geq \frac{1}{2}(r(k, \ell) + r(k, \ell+2)$ $(0 \leq \ell < N-1)$;

(4) $r(k+1, \ell) \geq \frac{1}{2}(r(k, \ell) + r(k+2, \ell)$ $(0 \leq k < N-1)$;

(5) $r(k, k)$ is divisible by k.

For $k = 0$ (resp. $\ell = 0$) we assume that $M_0(k, \ell)$ is the empty matrix having ℓ columns (resp. k rows), and $r(k, \ell) = 0$.

Proof. Initially put $M(k, \ell) = M_0(k, \ell)$ for every pair (k, ℓ). For a $k \times \ell$ matrix T let T^+ denote the $(k+1) \times \ell$ matrix obtained form T by appending a zero ($(k+1)$st) row, T^{++} is obtained by appending two zero rows. For $M = \sum_{i=1}^m B_i \otimes T_i$ we use M^+ for $\sum_{i=1}^m B_i \otimes T_i^+$, while $M^{++} = \sum_{i=1}^m B_i \otimes T_i^{++}$.

Let (k, ℓ) be a pair such that any of (1)–(5) is violated. Then we will replace some of the matrices $M(k', \ell')$ with matrices having larger rank. Over an infinite base field like $\mathbb{Q}$, each such replacement step (or each small group consisting of a few them) can be followed by an application of the data reduction procedure from [dGIR96] to keep intermediate (as well as the final) data small.

If (1) is violated then, like in [DM15], replace $M(k+1, \ell)$ with $M(k, \ell)^+$. We can treat a violation of (2) symmetrically.

When (3) is violated we consider the matrix $A = A(t) = M(k+2, \ell) + tM(k, \ell)^{++}$ as a $(k+2) \times \ell$ block matrix consisting of square blocks of size n from $\mathcal{B}$. We can choose t from any subset S of size $2nN + 1$ of the base field so that A has rank at least $r(k+2, \ell)$, while the first kn rows form a matrix of rank at least $r(k, \ell)$. This is because a necessary condition for violating either of these two conditions is that the determinant of an appropriate (but unknown) sub-matrix vanishes which determinant is, as a polynomial of degree at most nN in t is not identically zero. The product of these polynomials has degree at most $2nN$ therefore it cannot have more that $2nN$ zeros.

If A has rank larger than $r(k+2, \ell)$ then we replace $M(k+2, \ell)$ with A. Otherwise, like in [DM15], let U be the span of the first kn rows of A, V be the span of the first $(k+1)n$ rows and W be the span of the first kn rows and the last n rows. Note that these collections rows correspond to matrices of the form $A_0 = \sum B_i \otimes T_i$, $A_1 = \sum B_i \otimes T_i'$ and $A_2 = \sum B_i \otimes T_i''$ where T_i are $k \times \ell$ matrices, while T_i' and T_i'' have $(k+1)$ rows and ℓ columns. As $U \leq V \cap W$ and the row space of A is $V + W$, we have $r(k, \ell) \leq \dim U \leq \dim(V \cap W) = \dim V + \dim W - \dim V + W = \dim V + \dim W - r(k+2, \ell)$. It follows that $\dim V + \dim W \geq r(k, \ell) + r(k+2, \ell)$, whence violation of (3) is only possible if either $\dim V$ or $\dim W$ is strictly larger than $\frac{1}{2}(r(k, \ell) + r(k+2, \ell))$. Then we replace $M(k+1, \ell)$ with A_1 or A_2, according to which one has larger rank. A violation of (4) is treated symmetrically.

When (5) is violated then we can apply Lemma 2.1.

As in each round when violation of (1),. . .,(4) or (5) occurs the rank of at least one of the matrices $M(k,\ell)$ is incremented, the total number of rounds for achieving (1)–(5) is at most $N^3 n$. □

And here is essentially Proposition 2.10 of [DM15]. We include a proof (which is almost literally the same as the proof in [DM15]) here for completeness. We note that this lemma deals only with the property of certain families of functions, without referring to matrices.

Lemma 4.2 ([DM15, Proposition 2.10]). *Assume that $N > n > 0$, $r : \{0,1,\ldots,N\}^2 \to \mathbb{Z}$ is a function with $0 \le r(k,\ell) \le \min(k,\ell)n$ for $k,\ell \in \{0,1,\ldots,N\}$ also satisfying (1)–(5) of Lemma 4.1. Suppose further that $r(1,1) > 1$, and there exists d such that $n \le d+1 \le N$ and $r(d+1,d+1) = n(d+1)$. Then, $r(d,d) = nd$ as well.*

Proof. By $r(d+1,d+1) = n(d+1)$, for $1 \le a < d+1$,

$$r(d+1,a) \ge \frac{(d+1-1)\cdot r(d+1,0) + a \cdot r(d+1,d+1)}{d+1} = an.$$

As by assumption $r(d+1,a) \le an$, we have $r(d+1,a) = an$. Similarly $r(a,d+1) = an$ for $1 \le a < d+1$.

Then we bound $r(1,d)$ as follows:

$$\begin{aligned} r(1,d) &\ge \frac{(d-1)\cdot r(1,d+1) + 1 \cdot r(1,1)}{d} \\ &\ge \frac{(d-1)n+2}{d} = n - \frac{n-2}{d} > n-1. \end{aligned}$$

Note that we use $r(1,1) > 1$ and $d \ge n-1$. Since $r(1,d) \in \mathbb{Z}$, $r(1,d) = n$.

We are ready to bound $r(d,d)$ then.

$$\begin{aligned} r(d,d) &\ge \frac{(d-1)\cdot r(d+1,d) + 1 \cdot r(1,d)}{d} \\ &= \frac{(d-1)dn+n}{d} = nd - n + \frac{n}{d}. \end{aligned}$$

From $d \ge n-1$ it is inferred easily that $-n + \frac{n}{d} > -d$. Therefore $nd - n + \frac{n}{d} > (n-1)d$. By (5) we conclude that $r(d,d) = nd$. □

We finally remark that, if we use Lemma 4.1 in the proof of Theorem 1.1, then n in the statement of the lemma will be $r+1$, N will be d', $M_0(d',d')$ is the nonsingular $(r+1)d' \times (r+1)d'$ block of A and $M_0(p,q)$ can be actually even the zero matrix for $(p,q) \ne (d',d')$. It will prepare matrices in several not necessarily square blow-ups, among others, most importantly, one in an (r,r)-blow-up with a similar content as described in the proof of Theorem 1.1.

Acknowledgements. We would like to thank the authors of [GGOW15] and of [DM15] for sharing their ideas with us and making us possible to read early versions of their manuscripts. Part of the work was done when Gábor and Youming were visiting the Centre for Quantum Technologies at the National University of Singapore. Research of the first author was also supported in part by the Hungarian National Research, Development and Innovation Office – NKFIH, Grants NK105645 and 115288. Youming's research was supported by the Australian Research Council DECRA DE150100720.

References

[Der01] Harm Derksen. Polynomial bounds for rings of invariants. *Proceedings of the American Mathematical Society*, 129(4):955–964, 2001.

[dGIR96] Willem A. de Graaf, Gábor Ivanyos, and Lajos Rónyai. Computing Cartan subalgebras of Lie algebras. *Applicable Algebra in Engineering, Communication and Computing*, 7(5):339–349, 1996.

[DM15] Harm Derksen and Visu Makam. Polynomial degree bounds for matrix semi-invariants. preprint ArXiv:1512.03393, 2015.

[FR04] M. Fortin and C. Reutenauer. Commutative/noncommutative rank of linear matrices and subspaces of matrices of low rank. *Séminaire Lotharingien de Combinatoire*, 52:B52f, 2004.

[GGOW15] Ankit Garg, Leonid Gurvits, Rafael Oliveira, and Avi Wigderson. A deterministic polynomial time algorithm for non-commutative rational identity testing. preprint ArXiv:1511.03730, 2015.

[Gur04] Leonid Gurvits. Classical complexity and quantum entanglement. *J. Comput. Syst. Sci.*, 69(3):448–484, 2004.

[IKS10] Gábor Ivanyos, Marek Karpinski, and Nitin Saxena. Deterministic polynomial time algorithms for matrix completion problems. *SIAM J. Comput.*, 39(8):3736–3751, 2010.

[IQS15] Gábor Ivanyos, Youming Qiao, and K. V. Subrahmanyam. Non-commutative Edmonds' problem and matrix semi-invariants. preprint arXiv:1508.00690, 2015.

[Ram54] K.G. Ramanathan. *Lectures on the Algebraic Theory of Fields*. Tata Institute of Fundamental Research, Bombay, 1954.